

	1970	1971	1972	1973	1974	1975	1976	1977	1978	1979	1980	1981	1982	1983	1984	1985	1986	1987	1988	1989	1990	1991	1992	1993	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032	2033	2034	2035	2036	2037	2038	2039	2040	2041	2042	2043	2044	2045	2046	2047	2048	2049	2050	2051	2052	2053	2054	2055	2056	2057	2058	2059	2060	2061	2062	2063	2064	2065	2066	2067	2068	2069	2070	2071	2072	2073	2074	2075	2076	2077	2078	2079	2080	2081	2082	2083	2084	2085	2086	2087	2088	2089	2090	2091	2092	2093	2094	2095	2096	2097	2098	2099	2100	2101	2102	2103	2104	2105	2106	2107	2108	2109	2110	2111	2112	2113	2114	2115	2116	2117	2118	2119	2120	2121	2122	2123	2124	2125	2126	2127	2128	2129	2130	2131	2132	2133	2134	2135	2136	2137	2138	2139	2140	2141	2142	2143	2144	2145	2146	2147	2148	2149	2150	2151	2152	2153	2154	2155	2156	2157	2158	2159	2160	2161	2162	2163	2164	2165	2166	2167	2168	2169	2170	2171	2172	2173	2174	2175	2176	2177	2178	2179	2180	2181	2182	2183	2184	2185	2186	2187	2188	2189	2190	2191	2192	2193	2194	2195	2196	2197	2198	2199	2200	2201	2202	2203	2204	2205	2206	2207	2208	2209	2210	2211	2212	2213	2214	2215	2216	2217	2218	2219	2220	2221	2222	2223	2224	2225	2226	2227	2228	2229	2230	2231	2232	2233	2234	2235	2236	2237	2238	2239	2240	2241	2242	2243	2244	2245	2246	2247	2248	2249	2250	2251	2252	2253	2254	2255	2256	2257	2258	2259	2260	2261	2262	2263	2264	2265	2266	2267	2268	2269	2270	2271	2272	2273	2274	2275	2276	2277	2278	2279	2280	2281	2282	2283	2284	2285	2286	2287	2288	2289	2290	2291	2292	2293	2294	2295	2296	2297	2298	2299	2300	2301	2302	2303	2304	2305	2306	2307	2308	2309	2310	2311	2312	2313	2314	2315	2316	2317	2318	2319	2320	2321	2322	2323	2324	2325	2326	2327	2328	2329	2330	2331	2332	2333	2334	2335	2336	2337	2338	2339	2340	2341	2342	2343	2344	2345	2346	2347	2348	2349	2350	2351	2352	2353	2354	2355	2356	2357	2358	2359	2360	2361	2362	2363	2364	2365	2366	2367	2368	2369	2370	2371	2372	2373	2374	2375	2376	2377	2378	2379	2380	2381	2382	2383	2384	2385	2386	2387	2388	2389	2390	2391	2392	2393	2394	2395	2396	2397	2398	2399	2400	2401	2402	2403	2404	2405	2406	2407	2408	2409	2410	2411	2412	2413	2414	2415	2416	2417	2418	2419	2420	2421	2422	2
--	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	---

FOR

**INVENTOR:**

Prepared by:

**ANTONELLI, TERRY, STOUT & KRAUS, LLP**  
Suite 1800  
1300 North Seventeenth St.  
Arlington, VA 22209  
Phone: (703) 312-6600  
Fax: (703) 312-6666

## **METHOD OF READING A REMOTE MEMORY**

### **BACKGROUND**

#### **Field of the Invention**

5           This invention relates generally to methods and apparatus for transferring data between network devices. In particular, the present invention relates to methods and apparatus for reading a remote memory across a network.

#### **Description of the Related Art**

10           Conventional servers typically have multiple input/output (I/O) controllers, each supporting at least one I/O device, and a memory mapped load/store configuration. In the example of Fig. 1, there are a plurality of central processing units CPU<sub>1</sub> - CPU<sub>n</sub>, a host system bus and a system memory SM connected to the host system bus via a memory controller MC. An I/O bridge connects the memory controller MC to an I/O  
15 bus such as, for example, a Peripheral Component Interconnect (PCI) synchronous bus as described in the latest version of "PCI Local Bus Specification, Revision 2.1" set forth by the PCI Special Interest Group (SIG) on June 1, 1995. Connected to the I/O bus are an I/O controller A (usually implemented as a slot-based adaptor card) for a hard disk (not shown), an I/O controller B (usually implemented as a slot-based adaptor  
20 card) for a CD-ROM drive (not shown) and a network interface controller (NIC).

Suppose, for example, that CPU<sub>1</sub> wishes to transfer data to the hard disk via I/O controller A as shown in Fig. 2. CPU<sub>1</sub> first stores the write command and its associated data within a block within the system memory SM. CPU<sub>1</sub> transfers a command to the I/O controller A via a path over the system bus, I/O bridge, and I/O

5 bus. This tells the I/O controller A that a new command has been issued. I/O controller card A must then read the data from system memory SM using a pointer, which is the value representing an address within the system memory SM where the data associated with the command can be found. (The pointer may be virtual or physical and the location of the data is not necessarily contiguous with the location of  
10 the command. Indeed, the data may be split, requiring a Scatter/Gather List (SGL) to describe the locations of the data.) To get the block of data from the system memory SM back to I/O controller A may require several separate fetches. The data is then subsequently written from the I/O controller A to the hard disk HD. The CPU<sub>1</sub> must always load the data and the I/O controller must always separately read the write  
15 command to know where the data is located and perform the fetches to obtain the data.

A similar load/store procedure occurs when a CPU reads a block of data from the hard disk, i.e., the I/O controller A would store the block of data in a block within the system memory SM, then pass an indication to the CPU that the read process from the hard disk HD has been finished, whereupon the CPU must separately access the block  
20 within the system memory SM to obtain the data.

This conventional load/store procedure (illustrated generally in Fig. 3) of sending a command with pointer (step 1), waiting for and receiving a request for data (step 2) and subsequently sending the data in response to the request (step 3) has substantial inherent latencies and delays. Even though the CPUs perform optimally, the performance of the server can still be less than optimum because the procedure is very inefficient. The data transfers slow down the entire system and many CPU cycles will pass before they are completed. Although, the PCI bus architecture provides the most common accepted method used to extend computer systems for add-on arrangements (e.g., expansion cards) with new disk memory storage capabilities, it has performance limitations and scales poorly in server architectures. Furthermore, a server may have a significant number of I/O devices which are of radically different types, store different kinds of data and/or vary from each other in the addressing sequence by which the data blocks containing the data are written and read out.

A data transfer from another device across a network is similarly made without direct reference to the system memory. A network interface controller (NIC) acts as the communications intermediary between the device and the network and passes data blocks to and from the network in the speed and manner required by the network. The data transfer between the devices over the network is virtualized into a pair of starting and ending points corresponding to the NIC for each of the devices. Other parts of the devices, such as the I/O controllers and memory controller which controls the writing

and reading of the transferred data blocks to and from the device memory, are not involved when the data is transferred between the NICs across the network.

Furthermore, although not shown in Figs. 1 and 2, transport and other protocols (e.g., TCP, IP) are implemented at various levels of firmware and software in the device to

5 control, distinguish or review the transferred data in order to render the transfer of data over the network more reliable. The multiplexing and demultiplexing processes are computationally expensive and a CPU must control the movement of the transfer data blocks into and out of the memory controller or I/O controller during the transfer of each data block. Also, an intermediate copy of the data must be made in the hardware  
10 of the memory controller or I/O controller and at other levels or layers, mode switches and context switches of the device.

### SUMMARY

The present invention is directed to the transfer of data between network  
15 devices. In an example embodiment, a method of reading data from a remote device transfers data directly from the remote memory of the remote device to the local memory of the local device. A message is sent from the local device to the remote device which includes a transport header indicating the message type of the message. The remote device processes the message to determine whether or not the transport  
20 header of the message identifies the message as a type of remote Direct Memory

Access (rDMA) read operation. If the message is that type of remote Direct Memory Access (rDMA) read operation, then the remote device performs a remote Direct Memory Access (rDMA) write operation to the local device in accordance with data elements included in the message.

5

### **BRIEF DESCRIPTION OF THE DRAWINGS**

The foregoing and a better understanding of the present invention will become apparent from the following detailed description of example embodiments and the claims when read in connection with the accompanying drawings, all forming a part of the disclosure of the invention. While the foregoing and following written and illustrated disclosure focuses on disclosing example embodiments of the invention, it should be clearly understood that the same is by way of illustration and example only and is not to be taken by way of limitation.

10

15

Fig. 1 is a generalized block diagram providing a basic illustration of a server architecture having a conventional memory mapped load/store configuration.

Fig. 2 is a block diagram of the conventional server architecture showing the transfer of data blocks between a CPU and a hard disk.

20

Fig. 3 is a chart illustrating the steps of a conventional load/store I/O operation in which an I/O data block is transferred from a device on a network.

Fig. 4 is a diagram illustrating a server architecture in which a method according to an example embodiment of this invention may be advantageously performed.

Fig. 5 shows the structure of the data elements in the transport protocol of a  
5 rDMA read message according to an example embodiment of the invention.

Fig. 6 illustrates the processing of rDMA read messages according to an example embodiment of the invention.

Fig. 7 is a generalized block diagram showing the data transfer processing in a prior art network.

10 Fig. 8 is a generalized block diagram showing the VI architecture with which an example embodiment of the present invention may be used.

Fig. 9 is a block diagram of the VI architectural model.

Fig. 10 is a block diagram of the virtual interface in the VI architecture model.

Fig. 11 illustrates the processing of rDMA according to an example  
15 embodiment of the invention implemented across a virtual interface.

### **DETAILED DESCRIPTION**

For the reasons discussed in the Description of the Related Art, it is desirable for a device NIC to be able to push data across into a shared buffer pool of another  
20 device and direct the data into and from the device's memory. An example

embodiment of the present invention seeks to decrease the inefficiencies of the transfer of I/O data to input/output devices in a server architecture, such as what occurs when an I/O data block is transferred to or from a mass storage I/O device such as a hard disk. In particular, PCI compliant I/O adapters cannot accomplish data transfers

5 without the multiple steps discussed above. A server is a type of computer system having an architecture or otherwise designed to be able to support multiple I/O devices and to transfer data with other computer systems at high speed. (Due to recent advances in the performance and flexibility of computer systems, many modern computers are servers under this definition.) Although many servers currently utilize  
10 PCI buses, the example embodiment of the invention set forths a data transfer where the transferee device has remote direct memory access (RDMA) to virtual addresses, thus enabling protected, target-managed data transfer.

The example embodiment discloses a method which optimizes the transfer of I/O data blocks to reduce latency when one device attempts to read data from the  
15 remote memory of a remote device into its local memory either for the execution of instructions or the storage of the data. For a disk drive, for example, the data block is the unit of addressing and data transfer. If the value of one byte is to be updated on a disk, then the data transfer would include a block of data (512 bytes, for example) that contains the byte of interest. The byte value in the copy of the block held in memory  
20 would be updated, and then that block would be transferred from memory to the drive,



overwriting the old block stored on the disk. However, the method according to the example embodiment is not limited in its application to disk drives or I/O storage devices. In particular, the method according to the example embodiment may be useful for transferring data among computers and other devices on a network since data

5 latency is generally critical in such environments. The data may be transferred in blocks of different sizes depending upon, for example, the target device, the transfer protocol (such as, for example, ethernet packets), etc.

One example application of the invention is in a network interface controller (NIC) or chipset incorporated in the I/O unit of a server device configured to operate in

10 conjunction with one or more processors, such as the Intel Pentium II Xeon™ or Pentium III Xeon™ processor. However, such an embodiment is but one possible example of the invention which may, of course, be applied in any computer having a processor and an input/output device and indeed in any server architecture where an improvement in writing and reading data blocks to or from an I/O device is desired for

15 whatever reason.

An implementation of the example embodiment of the invention is in a host server architecture in a system area network (SAN) with a switched fabric configuration. There are several commercially available proprietary SAN fabrics, such as GigaNet and ServerNet by Compaq, however they do not support rDMA read

20 methods. Therefore, the example embodiment uses a Next Generation I/O (NGIO)

architecture allowing remote direct memory access in compliance with the NGIO standard, Version 1.0, published July 22, 1999. The NGIO architecture includes a fabric cluster based networking medium with new link specifications for a network of servers. It utilizes remoted I/O devices out on a System Area Network (SAN) fabric rather than I/O controllers and devices on system level I/O buses, such as PCI. NGIO is also more reliable because the I/O and server domains are separate and the CPU domains the I/O. Fault domains are separate so I/O devices going down will not necessarily affect the CPU infrastructure and memory buses. (This application of the method and apparatus of the invention is exemplary only. The method and apparatus of the invention are also applicable to non-switched point-to-point connection links in other configurations or networks such as System I/O.)

The example embodiment of a server architecture might be as shown in Fig. 4. Of course, a server can operate as a host as shown when initiating a data transfer or it can operate as a target device when receiving a data transfer initiated by another device. It is shown as a host having a plurality of CPUs such as an Intel Pentium II Xeon™ or Pentium III XEON™ processor, a host system bus, a memory controller, system memory, and an I/O unit containing a host channel adapter (HCA). The CPUs and memory controller of the host may be connected along the host system bus as shown in Fig. 4. The external data transfers are optimized through the host channel adapter (HCA) of the I/O unit at all times. The data transfer methods are available in

A NGIO. They may ~~be~~ also be available in other suitable architectures which are currently existing or later become available, such as System I/O. The commands and I/O data move through the host independently of the CPU. This helps the CPU or other elements of the host avoid having to expend system resources to accomplish

5 transfer of I/O data blocks since there may be access conflicts with other functions.

A The host channel adapter ~~CA~~ of the I/O unit, in turn, is connected to a switching fabric SF, which may contain many different switches SW. The switches are preferably multi-stage switches with naturally redundant communication channel links through the fabric such that a plurality of messages can be traveling through the

10 switching fabric at any given time. Each channel link between the HCA and the switches includes a request/response protocol permitting message sends, rDMA read and write, management and retry on transient errors. The channel links may be grouped together for additional bandwidth. On the other side of the link from the host channel adapter, a target channel adapter (TCA) adheres to all link semantics. The

15 TCA is a subset of a HCA and utilizes an appropriate interface (e.g., Gigabit Ethernet, SCSI, FC-AL, etc) to the I/O controller in the target device.

Accordingly, when a CPU issues a read command, for example, it simply passes the command to the host channel adaptor which injects it into the switched fabric SF, such that the CPU does not have to wait for processing of the and locking of the

20 system bus, but instead goes on to perform other processing operation until the

processing is completed. According to the present invention, the channel link is any means of transferring data, including but not limited to virtual channels, used to transfer data between two endpoints.

A

While the example embodiment is an NGIO implementation and thus supports

5 the channel link definition provided in the specification identified above, the present invention is not so limited. In accordance with the <sup>implementation</sup>~~implementation~~ in the NGIO specification or similar systems, once injected into the switched fabric SF, the write command travels through the switches and eventually arrives at a target channel adapter TCA where it can be given to an I/O controller where it is subsequently written to the

10 hard disk HD or to a network interface where it is subsequently transferred to another computer device on a connected network (not shown). Accordingly, the inherent delays in deciphering the command and writing of the data as required by the I/O controller are not experienced by the processor P which is on the other side of the switching fabric, and can continue processing. As shown in Fig. 4, the message

15 passing configuration can contain a plurality of channel adapters CA and host servers connected to the switching fabric SF, such that there can be a multitude of different messages traveling through the fabric and where all of the hosts can continue operating while their messages are traveling through the switching fabric. The example embodiment of the invention described in this application is implemented in a hardware

20 device compliant with the Next Generation I/O (NGIO) architecture. The example

embodiment and other embodiments of the invention can be implemented in conjunction with other types of switch fabric-based I/O architectures. The example embodiment NGIO uses a similar model for input/output data transfer as is specified by the VI architecture. A highly efficient interface such as the kernel interface in the  
5 example embodiment may thus be beneficially used by various computer devices having NGIO hardware connected to a network fabric. However, the example embodiment and other embodiments of the invention may also be used with non-NGIO hardware.

Conventional network interface controllers do not support rDMA read methods. The exemplary embodiment of the invention therefore provides a new method of  
10 rDMA read request processing. This rDMA read technique provides an I/O communication mechanism that allows a remote attached I/O device to read remote host memory with minimal host-side CPU involvement. It employs a simple transport protocol between the host-side requesting I/O data transfers and the device-side performing the I/O data transfer operations. The method provides for remote Direct  
15 Memory Access (rDMA) write and immediate data services to be utilized by the host-side to provide the necessary I/O data transfer and completion notification. The remote read processing is performed with new software techniques that are described in detail below.

To initiate a rDMA read request, a device-side driver provides a set of source  
20 buffers, which reference the remote host-side memory, and a set of destination buffers

that reference the local device-side memory. These buffers are pre-registered with their respective network interface controller (NIC) before the read request is initiated. The remote read mechanism allows these buffers to be specified as multiple data segments with offsets and a channel designation, which indicates the data path. The rDMA

- 5 Sequence field is used to specify the last data segment and completion of the rDMA read request.

The rDMA read request interface describes the data transfer from the host source buffers to the device destination buffers by using rDMA elements. The rDMA elements are embedded in a lightweight transport message that is identified by a unique header type and is recognized by the host-side as a rDMA read request. These rDMA elements along with the transport header, which specify both the type (TpInfoType) and count (TpInfoSize), are illustrated by Fig. 5. After receiving the rDMA read request, the host-side will schedule rDMA write transactions that are directed back to the device-side in order to complete the request. This rDMA read request is illustrated by Fig. 6.

Two separate initiated rDMA read requests are shown in Fig. 6 to help illustrate the design of the new rDMA read method. This example highlights the use of multiple segments to support scatter/gather and the use of multiple data channels that multiplex separate rDMA read transfers. The first rDMA read request (step 1) is scheduled for Data Channel 1 and the second rDMA read request (step 2) is schedule for Data

Channel 2 to avoid blocking and optimize performance. The rDMA read mechanism identifies each rDMA read with a unique sequence that will be used to complete the read request. The rDMA mechanism on the host-side recognizes the incoming packets on the control channel as embedded rDMA read requests and builds the proper rDMA write descriptors with the proper sequence inserted into the immediate data field on the last segment of each rDMA read request. The rDMA read completion for both requests are processed on the device-side, based on the immediate data that arrives with the last data segments (steps 3 and 4) of each rDMA transfer. This immediate data eliminates the need for the host-side to send a data transfer acknowledgement packet back to the device-side.

The example embodiment provides the capability for multiple data segments and strict in-order delivery across each channel. The multiple data segments give I/O devices an efficient mechanism to support a typical scatter/gather list of data buffers. The in-order scheduling of the rDMA read across each channel provides a guarantee of data transfer ordering. There is no need to acknowledge the completion of the rDMA write in the exemplary embodiment since it utilizes immediate data services. The rDMA write completion is identified with a sequence and end of data context. Thus, the receiving I/O unit can immediately identify the completing rDMA read request without waiting for an acknowledgement packet from the host. This method eliminates an acknowledgement packet and also provides a very fast completion mechanism on the

I/O unit of the receiving I/O device. Accordingly, a substantial portion of the previous inherent delays are avoided, thus achieving the advantages of the present invention.

The example embodiment allows multiplexed rDMA read transactions to occur across multiple channels to provide optimum performance and eliminate possible


5 blocking conditions. This method gives the I/O device the flexibility to separate the large and small data transfers so that small data transfers will not be blocked behind large transfers. As shown in Fig. 6, there is one control channel and two data channels. Alternatively, a single channel may be used to transmit the rDMA read message and the data across the channel adapters. Such arrangement is advantageous in  
10 that there is only a single channel, but is disadvantageous in that if a huge block of data (e.g., 2 gigabytes) is to be transmitted, such huge data block transmission will tie up the single channel, thus blocking all other messaging which can continue to operate and transmit messages across the fabric. Dual channels are generally more advantageous in that bulk data or left-over data transmission will never tie up the messaging channel.

15 There may also be a plurality of different data channels. With such arrangement, left-over data of a rDMA write operation can be sent over another one of the data channels. Such is advantageous in that only one data channel will be tied up transmitting any bulk or left-over data. In contrast, multiple channels can be used in parallel for faster transmission of the bulk or left-over data. More specifically, the



bulk or left-over data can be striped into equally sized blocks and transmitted over various channels.

There are different possible points in the server architecture to implement the method of the present invention. The first possible implementation is in software at the

5 NIC of the I/O device. This implementation supports the currently available NICs which do not support the NGIO standard by allowing an NGIO communication link to be emulated in software. An example implementation of the invention uses such NICs because they are freely available and less expensive than NGIO hardware. Preferably, the method is implemented as firmware, and accelerated with hardware support. A  
10 network may of course have I/O devices implementing <sup>Ins.</sup> ~~A<sub>1</sub>~~ 

Optimal rDMA read capabilities requires a well-defined communication interface. Therefore, another exemplary embodiment of the invention utilizes the communication interface defined by the Virtual Interface (VI) Architecture.

Specifically, a preferred implementation of the example embodiment is in the NIC or  
15 other input/output control circuit of a server or other network device having VI enabled hardware, as well as a processor such as the Intel Pentium II Xeon™ or Pentium III Xeon™ processor. However, such an implementation is but one possible application of the example embodiment which may, of course, be applied in any device having a processor and an input/output circuit and indeed in any server architecture where an

improvement in the speed of communications with a network is desired for whatever reason.

The centralized data transfer processing of the prior art (see Fig. 7) prohibits data transfers from realizing the potential raw hardware performance offered by underlying high-speed networks. The Virtual Interface (VI) Architecture (see Fig. 8) enables distributed applications to perform low overhead communication over System Area Networks (SANs), which are high-bandwidth, low-latency networks interconnecting nodes within a distributed computer system. VI is described in the Virtual Interface Architecture Specification, Version 1.0, December 16, 1997, jointly authored by Intel Corporation, Microsoft Corporation, and Compaq Computer Corporation, and makes it possible to perform low overhead communication using off-the shelf SAN hardware. The Virtual Interface architecture specification defines a standard interface to a SAN controller such that user applications and kernel applications can move data out onto a SAN with minimal impact on the CPU. It is designated for use in networks, such as SANs, which have very low error rates and high reliability levels. Transport errors are rare and considered catastrophic. The connection is broken when they occur.

Before the example embodiment of the present invention will be described in detail, the basic VI Architectural model will first be described with reference to Fig. 9. As will become apparent, the VI architecture provides direct data access to the network interface hardware in a fully protected fashion. It thus avoids intermediate data copies and

bypasses the operating system in order to achieve low latency, high bandwidth, data transfer.

As shown in Fig. 9, the VI Architectural model includes several different Virtual Interfaces (VIs) 20-1 to 20-3 between a VI user agent 14 and a VI network interface controller (NIC) hardware 18 (such as the GNN1000 of the cLAN product family, available from GigaNet, Inc.). The VI user agent 14 provides communications for software processes through a VI. The software processes may include general application programs 10 using legacy interfaces 12 (for example, an operating system communications facility such as, e.g., Sockets, Remote Procedure Call or RPC, MPI) between the general application programs 10 and VI user agent 14. The software process may also be application programs 11, such as database applications in the example of Fig. 9, which do not need legacy interfaces 12.

Each one of the VIs 20-1 to 20-3 in Fig. 9 can have the details shown in Fig. 10. (Only VI 20-1 is shown for the sake of simplicity.) VI 20-1 is configured to directly access the process' memory for data transfer operations between the software application and the network. In particular, it includes a pair of work queues, one for send operations (a send queue 19) and one for receive operations (receive queue 21). The work queues store one or more descriptors 23 between the time it is Posted (placed in the queue) and the time it is Done (when the VI NIC 18 has completed processing it). The descriptor 23 is a data structure recognizable by the VI NIC 18 that describes a data movement request, and it includes a list of segments (a control segment, an

optional address segment and one or more data segments). The control segment identifies the type of VI NIC 18 data movement operation to be performed and the status of a completed NIC data movement operation. The data segment describes a communications buffer for a VI NIC <sup>18 data</sup> ~~data~~ 18 movement operations. A receive queue

5 21 contains descriptors that describe where to place incoming data. A send queue 19 contains descriptors that describe the data to be transmitted. A pair of VIs are associated using connection primitives (e.g., VipConnectWait, VipConnectAccept, VipConnectRequest) to allow packets sent at one VI to be received at the other VI. A send doorbell 25 is provided for allowing the user software process or application to

10 notify the VI NIC 18 that work (a descriptor describing a requested data transfer operation) has been placed in send queue 19 and to request that the VI NIC 18 start the data transfer operations described by the descriptor placed in send queue 19. On the other hand, the receive doorbell 27 is not used to initiate data transfers. It is used simply to supply buffers for incoming data, such as message sends or immediate data

15 provided with inbound rDMA transfers, according to descriptors in receive queue 21.

Referring to Fig. 9 again, the VI user agent 14 is a software component that enables legacy interfaces 12 and database applications 11 to utilize VI NIC 18. The VI user agent 14, in part, abstracts the details of the underlying VI NIC 18 hardware in accordance with API(s) defined by the legacy interfaces 12. The APIs in VI user agent

20 14 primarily focus on providing primitives for synchronous communication with the VI

hardware; the asynchronous operations being facilitated by means of callbacks as discussed in the Description of the Related Art above. The VI user agent 14 includes a library of primitives, known as the Virtual Interface Primitives Library (VIPL) which receives commands from legacy interfaces 12 and provides one or more functions

5 corresponding to each command. The functions include those for creating a VI (VipCreateVI), for destroying a VI (VipDestroyVI), for connecting one VI to another VI (e.g., VipConnectWait, VipConnectRequest), for accepting or rejecting a VI connection request (VipConnectAccept or VipConnectReject), for terminating or disconnecting a connection between two VIs (VipDisconnect), to allow a process to register process  
10 memory with a VI NIC (VipRegisterMem), to post descriptors (to place a descriptor in a VI work queue using, e.g., VipPostSend, VipPostRecv), etc. Existing APIs to VI are targeted to user level applications. Details of the VI primitives (VIPL) are set forth in the VI Architecture Specification, version 1.0, December, 16, 1997 and the Virtual Interface (VI) Architecture Developer's Guide, Revision 1.0, September 9, 1998.

15 The VI kernel agent 16 is the privileged part of the device's operating system, usually a driver supplied by the vendor of VI NIC, that performs the setup and resource management functions needed to maintain a virtual interface. These functions include the creation/destruction of VIs, VI connection setup/teardown, interrupt management and or processing, management of system memory used by the VI NIC 18 and error handling. In  
20 the model, VI kernel agent 16 is accessed using the standard operating system mechanisms such as system calls. As shown by arrow 26 (Fig. 9), the legacy interfaces 12

make system calls to the VI kernel agent 16 to perform several control operations, including to create a VI on the local system, to connect the local VI to a VI on a remote system, and to register memory. The user software process or application must register memory to be used for data transfer prior to submitting the request for data transfer. The

5 memory regions used by descriptors and data buffers are registered prior to data transfer operations. Memory registration gives a VI NIC 18 a method to translate a virtual address to a physical address. The user receives an opaque memory handle as a result of memory registration. This allows a user to refer to a memory region using a memory

10 handle/virtual address pair without worrying about crossing page boundaries and keeping track of the virtual address to tag mapping. Memory registration enables the VI NIC 18 to transfer data directly between the registered buffers and the network. Traditional network transports often copy data between user buffers and one or more intermediate kernel buffers. Thus, processing overhead is decreased in the VI Architecture because data

15 ~~transfers between the~~ <sup>between the</sup> user software process or application and VI NIC 18 do not rely on system calls to the kernel.

After creating a VI on the local system, connecting the local VI to a remote VI and registering memory, legacy interfaces 12 or database applications 11 can use data transfer primitives of the VIPL library of VI user agent 14 to send and receive data. The VI Architecture defines two types of data transfer operations: 1) traditional send/receive

20 operations, and 2) Remote-DMA (RDMA) read/write operations. Once a connection is established, the legacy interfaces 12 or database applications 11 post the application's send

and receive requests directly to the local VI (to the send and receive queues 19 and 21).

They then post descriptors (e.g., place the descriptors in a work queue) and ring a doorbell to notify the VI NIC 18 that work has been placed in the work queue. The doorbell can be rung (notify the VI NIC 18 of the work in the queue) without kernel

5 processing. The VI NIC 18 then processes the descriptor by sending or receiving data, and then notifies the legacy interface 12 or database application 11 of the completed work using the completion queue 22. The processing of descriptors posted on a VI is performed in FIFO order but there is no implicit relationship between the processing of descriptors posted on different VIs. The VI NIC 18 directly performs the data transfer  
10 functions in response to the posted descriptors. This process removes the strict ordering between software layers in conventional network devices so that once the communication channel is established, the data is transferred more quickly.

The VI Architecture supports three levels of communication reliability at the NIC level: unreliable delivery, reliable delivery, and reliable reception. With reliable delivery  
15 and reception, corrupt data is detected, data is delivered exactly once, data order is guaranteed, data loss is detected and the connection is broken upon detection of an error. While providing high reliability, the VI architecture does not perform other transport level functions, including flow control, buffer management and fragmentation and reassembly. The VI Architecture Specification, version 1.0, December 16, 1997 states at page 15 that  
20 "VI consumers are responsible for managing flow control on a connection." The transport service provider of the present invention is designed to provide some transport level

functionality over the VI architecture or over an architecture similar to the VI architecture without adding unnecessary overhead.

As indicated previously, the rDMA read capability is optional in the VI specification. Therefore, this exemplary embodiment of the invention is appropriate for use as a new method of rDMA read request processing in VI-enabled systems. It employs the same simple transport protocol between the host-side requesting I/O data transfers and the device-side performing the I/O data transfer operations described above. In particular, it uses the Virtual Interface's remote Direct Memory Access (rDMA) write and immediate data services on the host-side to provide the necessary I/O data transfer and completion notification. Again, the rDMA read request interface describes the data transfer from the host source buffers to the device destination buffers by using rDMA elements. The rDMA elements are embedded in a lightweight transport message as shown in Fig. 5.

The rDMA read request in this exemplary embodiment is illustrated by Fig. 11 which shows certain elements of the VI architecture used in the method with reference to Figs. 9 and 10. To initiate a rDMA read request, a device-side I/O controller provides a set of source buffers in VI user agent 14, which reference the remote host-side memory, and a set of destination buffers in the I/O controller that reference the local device-side memory via send queues 19. These buffers are pre-registered with their respective VI network interface controller (NIC) 18 before the read request is



initiated. The remote read method allows these buffers to be specified as multiple data segments with offsets and a VI channel designation, which indicates the data path. The rDMA Sequence field is used to specify the last data segment and completion of the rDMA read request. After receiving the rDMA read request, the VI user agent 14 in the host-side will schedule rDMA write transactions that are directed back to the device-side via send queues 19 of the host and receive queues 21 of the I/O device in order to complete the request.

Two separate initiated rDMA read requests are shown in Fig. 11 to help illustrate the design of the new rDMA read method. This example highlights the use of multiple segments to support scatter/gather and the use of multiple data channels that multiplex separate rDMA read transfers. The first rDMA read request (step 1) is scheduled for a first send queue 19 and Data Channel 1 and the second rDMA read request (step 2) is scheduled for a second send queue 19 and Data Channel 2 to avoid blocking and optimize performance. The rDMA read mechanism identifies each rDMA read with a unique sequence that will be used to complete the read request. The VI user agent 14 on the host-side recognizes the incoming packets on the control channel as embedded rDMA read requests and builds the proper VI rDMA write descriptors with the proper sequence inserted into the immediate data field on the last segment of each rDMA read request. The rDMA read completion for both requests are processed on the device-side, based on the immediate data that arrives with the last data segments

(steps 3 and 4) of each rDMA transfer via receive queues 21. This immediate data eliminates the need for the host-side to send a data transfer acknowledgement packet back to the device-side.

This rDMA read technique provides an I/O communication mechanism that  
5 allows a remote attached I/O device to read remote host memory with minimal host-side CPU involvement using a virtual interface. It allows several multiplexed rDMA read transactions to occur across multiple VI channels to provide optimum performance and eliminate possible blocking conditions. This method gives the I/O device the flexibility to separate the large and small data transfers so that small data transfers will  
10 not be blocked behind large transfers.

The example embodiment provides the capability for multiple data segments and strict in-order delivery across each VI channel using send queues 19 and receive queues 21. The multiple data segments give I/O devices an efficient mechanism to support a typical scatter/gather list of data buffers. The in-order scheduling of the rDMA read  
15 across each channel provides a guarantee of data transfer ordering. There is no need to acknowledge the completion of the rDMA write in the exemplary embodiment since it utilizes the immediate data services provided by VI. The rDMA write completion is identified with a sequence and end of data context. Thus, the receiving I/O unit can immediately identify the completing rDMA read request without waiting for an  
20 acknowledgement packet from the host. This method eliminates an acknowledgement

packet and also provides a very fast completion mechanism on the I/O unit of the receiving I/O device. Accordingly, a substantial portion of the previous inherent delays are avoided, thus achieving the advantages of the present invention.

As shown in Fig. 11, there is one control channel and two data channels.

- 5 Alternatively, a single channel may be used to transmit the rDMA read message and the data across the channel adapters. Such arrangement is advantageous in that there is only a single channel, but is disadvantageous in that if a huge block of data (e.g., 2 gigabytes) is to be transmitted, such huge data block transmission will tie up the single channel, thus blocking all other messaging which can continue to operate and transmit
- 10 messages across the fabric. Dual channels are generally more advantageous in that bulk data or left-over data transmission will never tie up the messaging channel.

- There may also be a plurality of different data channels. With such arrangement, left-over data of a rDMA write operation can be sent over another one of the data channels. Such is advantageous in that only one data channel will be tied up
- 15 transmitting any bulk or left-over data. In contrast, multiple channels can be used in parallel for faster transmission of the bulk or left-over data. More specifically, the bulk or left-over data can be striped into equally sized blocks and transmitted over various channels.

- There are different possible points in the server architecture to implement the
- 20 method of the present invention. The first possible implementation is in software at the

VI NICs 18 of the host and the I/O device. This implementation supports the currently available NICs which do not support the NGIO standard or System I/O by allowing an NGIO communication link to be emulated in software. An example implementation of the invention uses such NICs because they are freely available and less expensive than

5 NGIO hardware. Preferably, the method is implemented as firmware, and accelerated with hardware support. A network may of course have I/O devices implementing different embodiments of the invention. Indeed, an advantage of the exemplary embodiment of the invention is that it is particularly useful and widely adaptable to any hardware in an I/O device having latency in data transfer operations.

10 Other features of the invention may be apparent to those skilled in the art from the detailed description of the example embodiments and claims when read in connection with the accompanying drawings. While the foregoing and following written and illustrated disclosure focuses on disclosing example embodiments of the invention, it should be understood that the same is by way of illustration and example

15 only, is not to be taken by way of limitation and may be modified in learned practice of the invention. While the foregoing has described what are considered to be example embodiments of the invention, it is understood that various modifications may be made therein and that the invention may be implemented in various forms and embodiments, and that it may be applied in numerous applications, only some of which have been

[illegible]